

PATENT
5500-54701
E0832CIP

"EXPRESS MAIL" MAILING LABEL NUMBER
EL 82201145405

DATE OF DEPOSIT April 2, 2001
I HEREBY CERTIFY THAT THIS PAPER OR
FEE IS BEING DEPOSITED WITH THE
UNITED STATES POSTAL SERVICE
"EXPRESS MAIL POST OFFICE TO
ADDRESSEE" SERVICE UNDER 37 C.F.R.
§1.10 ON THE DATE INDICATED ABOVE
AND IS ADDRESSED TO THE ASSISTANT
COMMISSIONER FOR PATENTS,
WASHINGTON, D.C. 20231.



Derrick Brown

Establishing an Operating Mode in a Processor

By:

Kevin J. McGrath
Michael T. Clark
James B. Keller

This application is a continuation-in-part of U.S. Patent Application Serial No. 09/483,101, filed January 14, 2000. This application claims benefit of priority to Provisional Application Serial No. 60/224,368 filed August 9, 2000.

5

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention is related to the field of processors and, more particularly, to
10 address and operand sizes in processors.

2. Description of the Related Art

The x86 architecture (also known as the IA-32 architecture) has enjoyed
15 widespread acceptance and success in the marketplace. Accordingly, it is advantageous to design processors according to the x86 architecture. Such processors may benefit from the large body of software written to the x86 architecture (since such processors may execute the software and thus computer systems employing the processors may enjoy increased acceptance in the market due to the large amount of available software).

20

As computer systems have continued to evolve, 64 bit address size (and sometimes operand size) has become desirable. A larger address size allows for programs having a larger memory footprint (the amount of memory occupied by the instructions in the program and the data operated upon by the program) to operate within
25 the memory space. A larger operand size allows for operating upon larger operands, or for more precision in operands. More powerful applications and/or operating systems may be possible using 64 bit address and/or operand sizes.

Unfortunately, the x86 architecture is limited to a maximum 32 bit operand size

and 32 bit address size. The operand size refers to the number of bits operated upon by the processor (e.g. the number of bits in a source or destination operand). The address size refers to the number of bits in an address generated by the processor. Thus, processors employing the x86 architecture may not serve the needs of applications which
5 may benefit from 64 bit address or operand sizes.

SUMMARY OF THE INVENTION

The problems outlined above are in large part solved by a processor as described
10 herein. The processor supports a first processing mode in which the address size is greater than 32 bits and the operand size may be 32 or 64 bits. The address size may be nominally indicated as 64 bits, although various embodiments of the processor may implement any address size which exceeds 32 bits, up to and including 64 bits, in the first processing mode. The first processing mode may be established by placing an enable
15 indication in a control register into an enabled state and by setting a first operating mode indication and a second operating mode indication in a segment descriptor to predefined states. Other combinations of the first operating mode indication and the second operating mode indication may be used to provide compatibility modes for 32 bit and 16 bit processing compatible with the x86 processor architecture (with the enable indication
20 remaining in the enabled state). Advantageously, 64 bit processing may be provided while providing compatibility with the x86 processor architecture, and hence supporting existing code written to the x86 processor architecture.

The processor may support the first processing mode in a variety of fashions. For
25 example, the processor may support the first processing mode directly in hardware. Alternatively, the processor may execute interpreter software which emulates instructions, including emulating the first processing mode. The processor may execute translation software which translates the instructions to native instructions executable by the processor. The native instructions may operate according to the first processing

mode. A combination of interpretation and translation may be used.

Broadly speaking, an apparatus is contemplated. The apparatus includes a first storage location, a second storage location, and a processor. The first storage location is
5 configured to store a segment selector identifying a segment descriptor including a first operating mode indication and a second operating mode indication. The second storage location is configured to store an enable indication. The enable indication, the first operating mode indication, and the second operating mode indication are indicative of an operating mode. The processor is configured to process an instruction according to the
10 operating mode.

Additionally, a method is contemplated. An operating mode is determined in response to an enable indication in a first storage location, a first operating mode indication in a segment descriptor, and a second operating mode indication in the segment
15 descriptor. Operands are fetched and addresses are generated in response to the operating mode.

BRIEF DESCRIPTION OF THE DRAWINGS

20 Other objects and advantages of the invention will become apparent upon reading the following detailed description and upon reference to the accompanying drawings in which:

Fig. 1 is a block diagram of one embodiment of a processor.

25

Fig. 2 is a block diagram of one embodiment of a segment descriptor for 32/64 mode.

Fig. 3 is a block diagram of one embodiment of a segment descriptor for

compatibility mode.

Fig. 4 is a block diagram of operation in compatibility mode and in legacy mode according to one embodiment of the processor shown in Fig. 1.

5

Fig. 5 is a table illustrating one embodiment of operating modes as a function of segment descriptor and control register values.

Fig. 6 is a table illustrating one embodiment of the use of instruction prefixes to
10 override default operating modes.

Fig. 7 is a block diagram of one embodiment of a register.

Fig. 8 is a diagram illustrating one embodiment of a global descriptor table and a
15 local descriptor table.

Fig. 9 is a block diagram of one embodiment of a 32/64 call gate descriptor.

Fig. 10 is a block diagram of an instruction format.

20

Fig. 11 is a block diagram of one embodiment of a computer system including the
processor shown in Fig. 1.

Fig. 12 is a block diagram of another embodiment of a computer system including
25 the processor shown in Fig. 1.

Fig. 13 is a block diagram of a second embodiment of a processor.

Fig. 14 is a block diagram of a prefix byte.

Fig. 15 is a flowchart illustrating one embodiment of an interpreter.

Fig. 16 is a flowchart illustrating one embodiment of a block from Fig. 15.

5

Fig. 17 is a flowchart illustrating one embodiment of another block from Fig. 15.

Fig. 18 is a flowchart illustrating one embodiment of a translator.

10 Fig. 19 is a block diagram illustrating one embodiment of mapping non-native architected state.

Fig. 20 is a block diagram illustrating a second embodiment of mapping non-native architected state.

15

Fig. 21 is a block diagram illustrating a third embodiment of mapping non-native architected state.

Fig. 22 is a block diagram of one embodiment of a carrier medium.

20

While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Turning now to Fig. 1, a block diagram illustrating one embodiment of a processor 10 is shown. Other embodiments are possible and contemplated. In the 5 embodiment of Fig. 1, processor 10 includes an instruction cache 12, an execution core 14, a data cache 16, an external interface unit 18, a memory management unit (MMU) 20, and a register file 22. In the illustrated embodiment, MMU 20 includes a set of segment registers 24, a first control register 26, a second control register 28, a local descriptor table register (LDTR) 30, and a global descriptor table register (GDTR) 32. Instruction 10 cache 12 is coupled to external interface unit 18, execution core 14, and MMU 20. Execution core 14 is further coupled to MMU 20, register file 22, and data cache 16. Data cache 16 is further coupled to MMU 20 and external interface unit 18. External interface unit 18 is further coupled to MMU 20 and to an external interface.

15 Processor 10 may employ a processor architecture compatible with the x86 architecture and including additional architectural features to support 64 bit processing. Processor 10 is configured to establish an operating mode in response to information stored in a code segment descriptor corresponding to the currently executing code and further in response to one or more enable indications stored in one or more control 20 registers. As used herein, an "operating mode" specifies default values for various programmably selectable processor attributes. For example, the operating mode may specify a default operand size and a default address size. The default operand size specifies the number of bits in an operand of an instruction, unless an instruction's encoding overrides the default. The default address size specifies the number of bits in an 25 address of a memory operand of an instruction, unless an instruction's encoding overrides the default. The default address size specifies the size of at least the virtual address of memory operands, and may also specify the size of the physical address. Alternatively, the size of the physical address may be independent of the default address size and may instead be dependent on the LME bit described below (e.g. the physical address may be

32 bits if the LME bit is clear and an implementation-dependent size greater than 32 bits and less than 64 bits if the LME bit is set) or on another control bit (e.g. the physical address extension bit, or PAE bit, in another control register). As used herein, a "virtual address" is an address generated prior to translation through an address translation

5 mechanism (e.g. a paging mechanism) to a "physical address", which is the address actually used to access a memory. Additionally, as used herein, a "segment descriptor" is a data structure created by software and used by the processor to define access control and status for a segment of memory. A "segment descriptor table" is a table in memory having multiple entries, each entry capable of storing a segment descriptor.

10

In the illustrated embodiment, MMU 20 generates an operating mode and conveys the operating mode to execution core 14. Execution core 14 executes instructions using the operating mode. More particularly, execution core 14 fetches operands having the default operand size from register file 22 or memory (through data cache 16, if the

15 memory operands are cacheable and hit therein, or through external interface unit 18 if the memory operands are noncacheable or miss data cache 16) unless a particular instruction's encoding overrides the default operand size, in which case the overriding operand size is used. Similarly, execution core 14 generates addresses of memory operands, wherein the addresses have the default address size unless a particular
20 instruction's encoding overrides the default address size, in which case the overriding address size is used. In other embodiments, the information used to generate the operating mode may be shadowed locally in the portions of processor 10 which use the operating mode (e.g. execution core 14), and the operating mode may be determined from the local shadow copies.

25

As mentioned above, MMU 20 generates the operating mode responsive to a code segment descriptor corresponding to the code being executed and further responsive to one or more values in control registers. Information from the code segment descriptor is stored in one of the segment registers 24 (a register referred to as CS, or code segment).

Additionally, control register 26 stores an enable indication (LME) which is used to enable an operating mode in which the default address size is greater than 32 bits ("32/64 mode") as well as certain compatibility modes for the 32 bit and 16 bit operating modes. The default operand size may be 32 bits in 32/64 mode, but instructions may override the 5 default 32 bit operand size with a 64 bit operand size when desired. If the LME indication is in an enabled state, then 32/64 mode may be used in addition to 32 bit and 16 bit modes. If the LME indication is in a disabled state, then 32/64 mode is disabled. In one embodiment, the default address size in 32/64 mode may be implementation-dependent but may be any value up to and including 64 bits. Furthermore, the size of the 10 virtual address may differ in a given implementation from the size of the physical address in that implementation.

It is noted that enable indications may be described herein as bits with the enabled state being the set state of the bit and the disabled state being the cleared state of the bit. 15 However, other encodings are possible, including encodings in which multiple bits are used and encodings in which the enabled state is the clear state and the disabled state is the set state. Accordingly, the remainder of this description may refer to the LME indication in control register 26 as the LME bit, with the enabled state being set and the disabled state being clear. However, other encodings of the LME indication are 20 contemplated, as set forth above.

Segment registers 24 store information from the segment descriptors currently being used by the code being executed by processor 10. As mentioned above, CS is one of segment registers 24 and specifies the code segment of memory. The code segment 25 stores the code being executed. Other segment registers may define various data segments (e.g. a stack data segment defined by the SS segment register, and up to four data segments defined by the DS, ES, FS, and GS segment registers). Fig. 1 illustrates the contents of an exemplary segment register 24A, including a selector field 24AA and a descriptor field 24AB. Selector field 24AA is loaded with a segment selector to activate

a particular segment in response to certain segment load instructions executed by execution core 14. The segment selector identifies the segment descriptor in a segment descriptor table in memory. More particularly, processor 10 may employ two segment descriptor tables: a local descriptor table and a global descriptor table. The base address 5 of the local descriptor table is stored in the LDTR 30. Similarly, the base address of the global descriptor table is stored in GDTR 32. A bit within the segment selector (the table indicator bit) selects the descriptor table, and an index within the segment selector is used as an index into the selected table. When an instruction loads a segment selector into one of segment registers 24, MMU 20 reads the corresponding segment descriptor from the 10 selected segment descriptor table and stores information from the segment descriptor into the segment descriptor field (e.g. segment descriptor field 24AB for segment register 24A). The information stored in the segment descriptor field may comprise any suitable subset of the segment descriptor, including all of the segment descriptor, if desired. Additionally, other information derived from the segment descriptor or other sources may 15 be stored in the segment descriptor field, if desired. For example, an embodiment may decode the operating mode indications from the code segment descriptor and store the decoded value rather than the original values of the operating mode indications. If an instruction causes CS to be loaded with a segment selector, the code segment may change and thus the operating mode of processor 10 may change. Segment descriptor tables are 20 described in more detail below.

In one embodiment, only the CS segment register is used in 32/64 mode. The data segment registers are ignored. In 16 and 32 bit modes, the code segment and data segments may be active. Furthermore, a second enable indication (PE) in control register 25 28 may affect the operation of MMU 20. The PE enable indication may be used to enable protected mode, in which segmentation and/or paging address translation mechanisms may be used. If the PE enable indication is in the disabled state, segmentation and paging mechanisms are disabled and processor 10 is in "real mode" (in which addresses generated by execution core 14 are physical addresses). Similar to the LME indication,

the PE indication may be a bit in which the enabled state is the bit being set and the disabled state is the bit being clear. However, other embodiments are contemplated as described above.

5 It is noted that MMU 20 may employ additional hardware mechanisms, as desired. For example, MMU 20 may include paging hardware to implement paging address translation from virtual addresses to physical addresses. The paging hardware may include a translation lookaside buffer (TLB) to store page translations.

10 It is noted that control registers 26 and 28 may be implemented as architected control registers (e.g. control register 26 may be CR4 and control register 28 may be CR0). Alternatively, one or both of the control registers may be implemented as model specific registers to allow for other uses of the architected control registers without interfering with 32/64 mode.

15 Generally, instruction cache 12 is a high speed cache memory for storing instruction bytes. Execution core 14 fetches instructions from instruction cache 12 for execution. Instruction cache 12 may employ any suitable cache organization, including direct-mapped, set associative, and fully associative configurations. If an instruction 20 fetch misses in instruction cache 12, instruction cache 12 may communicate with external interface unit 18 to fill the missing cache line into instruction cache 12. Additionally, instruction cache 12 may communicate with MMU 20 to receive physical address translations for virtual addresses fetched from instruction cache 12.

25 Execution core 14 executes the instructions fetched from instruction cache 12. Execution core 14 fetches register operands from register file 22 and updates destination registers in register file 22. The size of the register operands is controlled by the operating mode and any overrides of the operating mode for a particular instruction. Similarly, execution core 14 fetches memory operands from data cache 16 and updates

destination memory locations in data cache 16, subject to the cacheability of the memory operands and hitting in data cache 16. The size of the memory operands is similarly controlled by the operating mode and any overrides of the operating mode for a particular instruction. Furthermore, the size of the addresses of the memory operands generated by 5 execution core 14 is controlled by the operating mode and any overrides of the operating mode for a particular instruction.

Execution core 14 may employ any suitable construction. For example, execution core 14 may be a superpipelined core, a superscalar core, or a combination thereof. 10 Execution core 14 may employ out of order speculative execution or in order execution, according to design choice. Additionally, embodiments of execution core 14 may employ any of the above constructions and may include microcoding, as desired.

Register file 22 may include 64 bit registers which may be accessed as 64 bit, 32 15 bit, 16 bit, or 8 bit registers as indicated by the operating mode of processor 10 and any overrides for a particular instruction. The register format for one embodiment is described below with respect to Fig. 7. The registers included in register file 22 may include the RAX, RBX, RCX, RDX, RDI, RSI, RSP, and RBP registers (which may be 64 bit versions of the EAX, EBX, ECX, EDX, EDI, ESI, ESP, and EBP registers defined 20 in the x86 processor architecture, respectively). Additionally, in one embodiment, register file 22 may include additional registers addressed using a register extension (REX) prefix byte, described in more detail below. Register file 22 may further include the RIP register, which may be a 64 bit version of the EIP register. Alternatively, execution core 14 may employ a form of register renaming in which any register within 25 register file 22 may be mapped to an architected register. The number of registers in register file 22 may be implementation dependent for such an embodiment.

Data cache 16 is a high speed cache memory configured to store data. Data cache 16 may employ any suitable cache organization, including direct-mapped, set associative,

and fully associative configurations. If a data fetch or update misses in data cache 16, data cache 16 may communicate with external interface unit 18 to fill the missing cache line into data cache 16. Additionally, if data cache 16 employs a writeback caching policy, updated cache lines which are being cast out of data cache 16 may be

5 communicated to external interface unit 18 to be written back to memory. Data cache 16 may communicate with MMU 20 to receive physical address translations for virtual addresses presented to data cache 16.

External interface unit 18 communicates with portions of the system external to
10 processor 10. External interface unit 18 may communicate cache lines for instruction cache 12 and data cache 16 as described above, and may communicate with MMU 20 as well. For example, external interface unit 18 may access the segment descriptor tables and/or paging tables on behalf of MMU 20.

15 It is noted that processor 10 may include an integrated level 2 (L2) cache, if desired. Furthermore, external interface unit 18 may be configured to communicate with a backside cache in addition to communicating with the system.

While the processor architecture described herein may be compatible with the x86
20 processor architecture for 16 and 32 bit modes, in one embodiment, other embodiments may employ any 16 and 32 bit modes. The other embodiments may or may not be compatible with the x86 processor architecture or any other processor architecture. It is further noted that, while a specific set of information is described herein as being used to generate the operating mode, any combination of indications and/or information from
25 memory data structures such as segment descriptor tables and page tables may be used to generate the operating mode in various embodiments.

Turning now to Fig. 2, a block diagram of one embodiment of a code segment descriptor 40 for 32/64 mode is shown. Other embodiments are possible and

contemplated. In the embodiment of Fig. 2, code segment descriptor 40 comprises 8 bytes with the most significant 4 bytes illustrated above the least significant 4 bytes. The most significant four bytes are stored at a numerically larger address than the least significant four bytes. The most significant bit of each group of four bytes is illustrated 5 as bit 31 in Fig. 2 (and Fig. 3 below), and the least significant bit is illustrated as bit 0. Short vertical lines within the four bytes delimit each bit, and the long vertical lines delimit a bit but also delimit a field (both in Fig. 2 and in Fig. 3).

Unlike the 32 bit and 16 bit code segment descriptors illustrated in Fig. 3 below, 10 code segment descriptor 40 does not include a base address or limit. Processor 10 employs a flat virtual address space for 32/64 mode (rather than the segmented linear address space employed in 32 bit and 16 bit modes). Accordingly, the portions of code segment descriptor 40 which would otherwise store the base address and limit are reserved in segment descriptor 40. It is noted that a virtual address provided through 15 segmentation may also be referred to herein as a "linear address". The term "virtual address" encompasses any address which is translated through a translation mechanism to a physical address actually used to address memory, including linear addresses and other virtual addresses generated in non-segmented architectures.

20 Segment descriptor 40 includes a D bit 42, an L bit 44 (set to one for a 32/64 mode code segment), an available bit (AVL) 46, a present (P) bit 48, a descriptor privilege level (DPL) 50, and a type field 52. D bit 42 and L bit 44 are used to determine the operating mode of processor 10, as illustrated in Fig. 5 below. AVL bit 46 is available for use by system software (e.g. the operating system). P bit 48 is used to 25 indicate whether or not the segment is present in memory. If P bit 48 is set, the segment is present and code may be fetched from the segment. If P bit 48 is clear, the segment is not present and an exception is generated to load the segment into memory (e.g. from disk storage or through a network connection). The DPL indicates the privilege level of the segment. Processor 10 employs four privilege levels (encoded as 0 through 3 in the DPL

field, with level 0 being the most privileged level). Certain instructions and processor resources (e.g. configuration and control registers) are only executable or accessible at the more privileged levels, and attempts to execute these instructions or access these resources at the lower privilege levels result in an exception. When information from code segment 40 is loaded into the CS segment register, the DPL becomes the current privilege level (CPL) of processor 10. Type field 52 encodes the type of segment. For code segments, the most significant bit two bits of type field 52 may be set (the most significant bit distinguishing a code or data segment from a system segment, and the second most significant bit distinguishing a code segment from a data segment), and the remaining bits may encode additional segment type information (e.g. execute only, execute and read, or execute and read only, conforming, and whether or not the code segment has been accessed).

It is noted that, while several indications in the code segment descriptor are described as bits, with set and clear values having defined meanings, other embodiments may employ the opposite encodings and may use multiple bits, as desired. Thus, for example, the D bit 42 and the L bit 44 may each be an example of an operating mode indication which may be one or more bits as desired, similar to the discussion of enable indications above.

Turning now to Fig. 3, a block diagram of one embodiment of a code segment descriptor 54 for 32 and 16 bit compatibility mode is shown. Other embodiments are possible and contemplated. As with the embodiment of Fig. 2, code segment descriptor 54 comprises 8 bytes with the most significant 4 bytes illustrated above the least significant 4 bytes.

Code segment descriptor 54 includes D bit 42, L bit 44, AVL bit 46, P bit 48, DPL 50, and type field 52 similar to the above description of code segment descriptor 40. Additionally, code segment descriptor 54 includes a base address field (reference

numerals 56A, 56B, and 56C), a limit field (reference numerals 57A and 57B) and a G bit 58. The base address field stores a base address which is added to the logical fetch address (stored in the RIP register) to form the linear address of an instruction, which may then optionally be translated to a physical address through a paging translation 5 mechanism. The limit field stores a segment limit which defines the size of the segment. Attempts to access a byte at a logical address greater than the segment limit are disallowed and cause an exception. G bit 58 determines the scaling of the segment limit field. If G bit 58 is set the limit is scaled to 4K byte pages (e.g. 12 least significant zeros are appended to the limit in the limit field). If G bit 58 is clear, the limit is used as is.

10

It is noted that code segment descriptors for 32 and 16 bit modes when 32/64 mode is not enabled via the LME bit in control register 26 may be similar to code segment descriptor 54, except the L bit is reserved and defined to be zero. It is further noted that, in 32 and 16 bit modes (both compatibility mode with the LME bit set and 15 modes with the LME bit clear) according to one embodiment, data segments are used as well. Data segment descriptors may be similar to code segment descriptor 54, except that the D bit 42 is defined to indicate the upper bound of the segment or to define the default stack size (for stack segments).

20

Turning next to Fig. 4, a diagram illustrating exemplary uses of the LME bit in control register 26 and the compatibility modes to allow for a high degree of flexibility in implementing the 32/64 mode and the 32 and 16 bit modes is shown. A box 60 illustrates exemplary operation when the LME bit is set, and a box 62 illustrates exemplary operation when the LME bit is clear.

25

As illustrated in box 60, the compatibility modes supported when the LME bit is set may allow for a 64 bit operating system (i.e. an operating system designed to take advantage of the virtual and physical address spaces in excess of 32 bits and/or data operands of 64 bits) to operate with a 32 bit application program (i.e. an application

program written using 32 bit operand and address sizes). The code segment for the operating system may be defined by the 32/64 mode code segment descriptor 40 illustrated in Fig. 2, and thus the L bit may be set. Accordingly, the operating system may take advantage of the expanded virtual address space and physical address space for the 5 operating system code and the data structures maintained by the operating system (including, e.g. the segment descriptor tables and the paging translation tables). The operating system may also use the 64 bit data type defined in 32/64 mode using instruction encodings which override the default 32 bit operand size. Furthermore, the operating system may launch a 32 bit application program by establishing one or more 32 10 bit compatibility mode segment descriptors (L bit cleared, D bit set, e.g. segment descriptor 54 shown in Fig. 2) in the segment descriptor table and branching into one of the compatibility mode segments. Similarly, the operating system may launch a 16 bit application program by establishing one or more 16 bit compatibility mode segment descriptors (L bit cleared, D bit cleared, e.g. segment descriptor 54 shown in Fig. 2) in the 15 segment descriptor table and branching into one of the compatibility mode segments. Accordingly, a 64 bit operating system may retain the ability to execute existing 32 bit and 16 bit application programs in the compatibility mode. A particular application program may be ported to 32/64 mode if the expanded capabilities are desired for that program, or may remain 32 bit or 16 bit.

20 While processor 10 is executing the 32 bit application program, the operating mode of processor 10 is 32 bit. Thus, the application program may generally execute in the same fashion as it does in 32 bit mode with the LME bit clear (e.g. when the operating system is a 32 bit operating system as well). However, the application program may call 25 an operating system service, experience an exception, or terminate. In each of these cases, processor 10 may return to executing operating system code (as illustrated by arrow 64 in Fig. 4). Since the operating system code operates in 32/64 mode, the address of the operating system service routine, exception handler, etc. may exceed 32 bits. Thus, processor 10 may need to generate an address greater than 32 bits prior to returning to the

operating system code. The LME bit provides processor 10 with an indication that the operating system may be operating in 32/64 mode even though the current operating mode is 32 bit, and thus processor 10 may provide the larger address space for operating system calls and exceptions.

5

In one embodiment, exceptions are handled using interrupt segment descriptors stored in an interrupt segment descriptor table. If the LME bit is set, the interrupt segment descriptors may be 16 byte entries which include a 64 bit address of the operating system routine which handles the exception. If the LME bit is clear, the 10 interrupt segment descriptors may be eight byte entries which include a 32 bit address. Accordingly, processor 10 accesses the interrupt descriptor table responsive to the LME indication (i.e. reading a 16 byte entry if the LME bit is set and reading an eight byte entry if the LME bit is clear). Therefore, exceptions may be handled by the 64 bit operating system even though the application program is executing in 32 bit compatibility mode. 15 Furthermore, processor 10 supports a 32 bit (or 16 bit) operating system if the LME bit is clear.

Similarly, the call mechanisms within processor 10 may operate in different fashions based on the state of the LME bit. Since the operating system typically executes 20 at a higher privilege level than the application program, transfers from the application program to the operating system are carefully controlled to ensure that the application program is only able to execute permitted operating system routines. More generally, changes in privilege level are carefully controlled. In one embodiment, processor 10 may support at least two mechanisms for performing operating system calls. One method may 25 be through a call gate in the segment descriptor tables (described in more detail below). Another method may be the SYSCALL instruction supported by processor 10, which uses a model specific register as the source of the address of the operating system routine. Updating the model specific registers is a privileged operation, and thus only code executing at a higher privilege level (e.g. operating system code) may establish the

address in the model specific register used by the SYSCALL instruction. For the SYSCALL method, a second model specific register may be defined to store the most significant 32 bits of the address of the operating system routine. Thus, if the LME bit is set, the address may be read from the two model specific registers. If the LME bit is 5 clear, the address may be read from the model specific register storing the least significant 32 bits. Alternatively, the model specific register used by the SYSCALL instruction may be expanded to 64 bits and the address may be 32 bits (the least significant 32 bits of the model specific register) or 64 bits based on the state of the LME bit.

10 As illustrated above, having the LME bit set may allow for processor 10 to operate in a system in which the operating system is 64 bit and one or more application programs are not 64 bit (e.g. 32 bit as shown or 16 bit, which operates in a similar fashion to the above description). Generally, even though the processor may be operating in 32 or 16 bit mode, the LME bit informs the processor that the operating system data structures 15 are as defined for the 64 bit modes, and the processor may access the structures appropriately. Additionally, as illustrated by box 62, having the LME bit clear may allow for processor 10 to operate in 32 bit or 16 bit modes compatible with the x86 architecture. As described above, the mechanisms for handling exceptions and operating system calls are designed to handle the LME bit being set or clear, and thus the 32 bit and 16 bit 20 modes may operate unmodified, even though processor 10 is capable of operating in 32/64 mode. Furthermore, by providing the x86 compatible 16 and 32 bit modes when the LME bit is clear, (and ignoring the L bit, which is reserved in these modes) processor 10 may operate in a system in which the L bit is defined for some other purpose than for 32/64 mode and may still support 32/64 mode if the LME bit is set. Accordingly, a 25 system employing a 32 bit operating system and 32 bit or 16 bit application programs may employ processor 10. Subsequently, the system could be upgraded to a 64 bit operating system without having to change processor 10.

Not illustrated in Fig. 4 is a 64 bit operating system and a 64 bit application

program operating with the LME bit set. The mechanisms for calling operating system routines described above for the 64 bit operating system and 32 bit application program may apply equally to the 64 bit application program as well. Additionally, call gates which support 64 bits of offset are supported (as will be described in more detail below).

5

Turning next to Fig. 5, a table 70 is shown illustrating the states of the LME bit, the L bit in the code segment descriptor, and the D bit in the code segment descriptor and the corresponding operating mode of processor 10 according to one embodiment of processor 10. Other embodiments are possible and contemplated. As table 70 illustrates,

- 10 if the LME bit is clear, then the L bit is reserved (and defined to be zero). However, processor 10 may treat the L bit as a don't care if the LME bit is clear. Thus, the x86 compatible 16 bit and 32 bit modes may be provided by processor 10 if the LME bit is clear. If the LME bit is set and the L bit in the code segment is clear, then a compatibility operating mode is established by processor 10 and the D bit selects 16 bit or
- 15 32 bit mode. If the LME bit and the L bit are set and the D bit is clear, 32/64 mode is selected for processor 10. Finally, the mode which would be selected if the LME, L and D bits are all set is reserved.

As mentioned above and illustrated in Fig. 6 below, the 32/64 operating mode

- 20 includes a default address size in excess of 32 bits (implementation dependent but up to 64 bits) and a default operand size of 32 bits. The default operand size of 32 bits may be overridden to 64 bits via a particular instruction's encoding. The default operand size of 32 bits is selected to minimize average instruction length (since overriding to 64 bits involves including an instruction prefix in the instruction encoding which may increase
- 25 the instruction length) for programs in which 32 bits are sufficient for many of the data manipulations performed by the program. For such programs (which may be a substantial number of the programs currently in existence), moving to a 64 bit operand size may actually reduce the execution performance achieved by the program (i.e. increased execution time). In part, this reduction may be attributable to the doubling in size in

memory of the data structures used by the program when 64 bit values are stored. If 32 bits is sufficient, these data structures would store 32 bit values. Thus, the number of bytes accessed when the data structure is accessed increases if 64 bit values are used where 32 bit values would be sufficient, and the increased memory bandwidth (and increased cache space occupied by each value) may cause increased execution time.

5 Accordingly, 32 bits is selected as the default operand size and the default may be overridden via the encoding of a particular instruction.

Turning next to Fig. 6, a table 72 is shown illustrating one embodiment of the use
10 of instruction prefixes to override the operating mode for a particular instruction. Other
embodiments are possible and contemplated. Execution core 14 determines the address
size and operand size for a particular instruction according to table 72. In particular for
the embodiment illustrated in Fig. 6, an instruction prefix byte (the address size override
prefix byte) may be used to override the default address size and another instruction
15 prefix byte (the operand size override prefix byte) may be used to override the default
operand size. The address size override prefix byte is encoded as 67 (in hexadecimal)
and the operand size override prefix byte is encoded as 66 (in hexadecimal). The number
of override prefixes in the particular instruction forms the columns of the table. The rows
20 of the table indicate the operand size and address size of the particular instruction, based
on the operating mode and the number of override prefixes in the corresponding column.
The number of override prefixes refers to the number of override prefixes of the
corresponding type (e.g. address size rows are the address size based on the number of
address size override prefixes and operand size rows are the operand size based on the
number of operand size override prefixes).

25

The column labeled "0" for the number of override prefixes illustrates the default
operand size and address size for each operating mode. It is noted that the 32 bit and 16
bit mode rows refer to both the compatibility modes (LME set) and the standard modes
(LME clear). Furthermore, while the default address size is 64 bits in 32/64 mode, the

actual number of address bits may be implementation dependent, as discussed above.

The inclusion of one address size override prefix in 32/64 bit mode changes the address size from 64 bit (which may be less than 64 bits for a given implementation but is greater than 32 bits) to 32 bit, as shown in table 72. Additionally, the inclusion of one operand size override prefix in 32/64 bit mode changes the operand size from 32 bit to 64 bit. It may be desirable to provide for a 16 bit operand as well (e.g. to support the short integer data type in the "C" programming language). Accordingly, the inclusion of two operand size override prefixes in 32/64 mode selects an operand size of 16 bits. The inclusion of more than two operand size override prefixes results in the same operand size as the inclusion of two operand size override prefixes. Similarly, the inclusion of more than one address size override prefix results in the same address size as the inclusion of one address size override prefix.

For the 32 bit modes, the inclusion of one override prefix toggles the default 32 bit size to 16 bit, and the inclusion of more than one override prefix has the same effect as the inclusion of one override prefix. Similarly, for 16 bit modes, the inclusion of one override prefix toggles the default 16 bit size to 32 bit, and the inclusion of more than one override prefix has the same effect as the inclusion of one override prefix.

Turning now to Fig. 7, a diagram illustrating one embodiment of the RAX register 74 is shown. Other registers within register file 22 may be similar. Other embodiments are possible and contemplated. In the embodiment of Fig. 7, register 74 includes 64 bits, with the most significant bit labeled as bit 63 and the least significant bit labeled as bit 0. Fig. 7 illustrates the portions of the RAX register accessed based upon the operand size of an instruction (if the A register is selected as an operand). More particularly, the entirety of register 74 is accessed if the operand size is 64 bits (as illustrated by the brace labeled "RAX" in Fig. 7). If the operand size is 32 bits, bits 31:0 of register 74 are accessed (as illustrated by the brace labeled "EAX" in Fig. 7). If the operand size is 16 bits, bits 16:0

of the register are accessed (as illustrated by the brace labeled "AX" in Fig. 7). The above operand sizes may be selected based on the operating mode and the inclusion of any override prefixes. However, certain instruction opcodes are defined which access an eight bit register (AH or AL in Fig. 7).

5

Turning next to Fig. 8, a block diagram is shown illustrating one embodiment of a global descriptor table 80 and a local descriptor table 82. Other embodiments are possible and contemplated. As illustrated in Fig. 8 and mentioned above, the base address of global descriptor table 80 is provided by GDTR 32 and the base address of local descriptor table 82 is provided by LDTR 30. Accordingly, to support placing global descriptor table 80 and local descriptor table 82 arbitrarily within the virtual address space, GDTR 32 and LDTR 30 may store 64 bit base addresses. If the LME bit is clear, the least significant 32 bits of the base address may be used to locate the descriptor tables.

15 Both global descriptor table 80 and local descriptor table 82 are configured to store segment descriptors of various types. For example, 32/64 mode code segment descriptors 84, 86, and 90 and compatibility mode descriptors 92 and 94 are illustrated in Fig. 8. Each of descriptors 84-94 occupies an entry in the corresponding descriptor table, where an entry is capable of storing one segment descriptor (e.g. 8 bytes for the 20 embodiments illustrated in Figs. 2 and 3). Another type of descriptor in global descriptor table 80 is a local descriptor table descriptor 96, which defines a system segment for the local descriptor table 82 and provides the base address stored in LDTR 30. LDTR 30 is initialized using an LLDT instruction having as an operand a segment selector locating descriptor 96 in global descriptor table 80. Global descriptor table 80 may store multiple 25 LDT descriptors locating different local descriptor tables, if desired. Since the LDT descriptor 96 may store a 64 bit offset if the LME bit is set, LDT descriptor 96 may occupy two entries in global descriptor table 80. If the LME bit is clear, LDT descriptor 96 may occupy a single entry in global descriptor table 80. Similarly, each task may have a task state segment (TSS) descriptor in one of descriptor tables 80 and 82 to store certain

information related to the task. Accordingly, a TSS descriptor may occupy two entries to allow for TSS information to be stored anywhere in the 64 bit address space.

The local and global descriptor tables may also store a call gate descriptor. For
5 example, Fig. 8 illustrates call gate descriptors 100, 102, and 104. Call gate descriptors support a 64 bit offset as well, and thus may occupy two entries in the corresponding descriptor table as well. An exemplary 32/64 call gate descriptor is illustrated in Fig. 9 below. Interrupt gate and trap gate descriptors may similarly occupy two entries in the corresponding descriptor table.
10

By maintaining the segment descriptor tables 80 and 82 at 8 bytes and using two entries for descriptors which include 64 bit offsets, descriptors for 16 and 32 bit modes may be stored in the same tables as the descriptors which include 64 bit offsets. Thus,
15 applications operating in compatibility modes may have appropriate descriptors in the same segment descriptor tables as the 64 bit operating systems.

Generally, call gates are used to manage the transition between a code segment having a lesser privilege level and a code segment have a greater privilege level (e.g. an application program calling an operating system routine). The lesser privileged code
20 includes a call or other branch instruction specifying, as a target, a segment selector (and an offset into the segment, which is ignored in this case). The segment selector identifies a call gate descriptor within the descriptor tables, which includes a minimum privilege level required to execute the greater privilege level code. When processor 10 executes the call or other branch instruction, processor 10 indexes the descriptor tables with the
25 segment selector and locates the call gate. If the current privilege level of processor 10 and the requestor privilege level (which is part of the segment selector, and may be used to lower the current privilege level for privilege checking purposes) both reflect sufficient privilege (e.g. the privilege levels are numerically less than or equal to the minimum privilege level in the call gate descriptor), then the call may proceed. The call gate

descriptor includes a segment selector for the target segment (the code segment having the greater privilege level) and the offset within the target segment at which code fetching is to begin. Processor 10 extracts the segment selector and the offset from the call gate descriptor and reads the target segment descriptor to begin fetching the code having the
5 greater privilege level. On the other hand, if either the current privilege level or the requestor privilege level is a lesser privilege level than the minimum privilege level in the call gate descriptor (e.g. either the current or requestor privilege level is numerically greater than the minimum privilege level), processor 10 signals an exception after accessing the call gate descriptor and without accessing the target descriptor. Thus,
10 access to code executing at greater privilege levels is carefully controlled.

As mentioned above, the call gate descriptor includes a target segment selector and offset within the segment. The reference to the target segment descriptor is illustrated in Fig. 8 as an arrow from a call gate descriptor to another descriptor. For
15 example, call gate descriptor 100 references mode descriptor 90; call gate descriptor 102 references 32/64 mode descriptor 86, and call gate descriptor 104 references 32/64 mode descriptor 84. As Fig. 8 illustrates, a call gate descriptor may be stored in either descriptor table and may reference a descriptor in the other table or in the same table.
Furthermore, a call gate descriptor may reference either a 32/64 mode descriptor or a
20 compatibility mode descriptor.

Generally, when processor 10 reads a descriptor from one of the descriptor tables using a segment selector, one descriptor table entry is read. However, if the LME bit is set and processor 10 detects that the entry is a call gate descriptor, an LDT descriptor, a
25 TSS descriptor, a trap gate descriptor, or an interrupt gate descriptor, processor 10 reads the next succeeding entry in the table to obtain the remainder of the descriptor. Accordingly, call gate descriptors, LDT descriptors, TSS descriptors, trap gate descriptors, and interrupt gate descriptors may coexist in a table with compatibility mode descriptors (or standard mode descriptors) which are of a different size, without

redefining the size of the table entries nor how the table is managed for descriptors which occupy one entry. Furthermore, since the second portion of the call gate descriptor, the LDT descriptor, the TSS descriptor, the trap gate descriptor, and the interrupt gate descriptor may be accessed as a segment descriptor, the portion of the descriptor which would be the type field of a descriptor in the second portion is set to an invalid type when the descriptor is stored into the descriptor table, as shown below in Fig. 9. Alternatively, processor 10 may read two consecutive entries from a descriptor table each time a descriptor table read is performed, and the second entry may be used if the first entry is a call gate, interrupt gate, trap gate, LDT descriptor type, or TSS descriptor type.

10

It is noted that code operating in any operating mode (32/64 mode, 32 bit compatibility mode, or 16 bit compatibility mode) may reference a call gate descriptor when the LME bit is set. Thus, a 32 or 16 bit application may call an operating system routine even if the address of the routine is outside the 32 bit or 16 bit address space using the call gate mechanism. Additionally, a call gate descriptor may reference a code segment having any operating mode. The operating system may ensure that the most significant 32 bits of the offset in the call gate are zero (for a 32 bit target segment) or the most significant 48 bits of the offset in the call gate are zero (for a 16 bit target segment).

20

Turning now to Fig. 9, a block diagram of one embodiment of a call gate descriptor 120 is shown. Other embodiments are possible and contemplated. Similar to Figs. 2 and 3, the most significant bytes are illustrated above the least significant bytes. The most significant bit of each group of four bytes is illustrated as bit 31 and the least significant bit is illustrated as bit 0. Short vertical lines within the four bytes delimit each bit, and the long vertical lines delimit a bit but also delimit a field. As mentioned above, a call gate descriptor occupies two entries in a descriptor table. The horizontal dashed line in Fig. 9 divides call gate descriptor 120 into an upper portion (above the line) and a lower portion (below the line). The lower portion is stored in the entry indexed by the call gate's segment selector, and the upper portion is stored in the next succeeding entry.

Call gate descriptor 120 includes a target segment selector (field 122), an offset (fields 124A, 124B, and 124C), a present (P) bit 126, a descriptor privilege level (DPL) 128, a type field 130, and a pseudo-type field 132. The P bit is similar to P bit 48 described above. The target segment selector identifies an entry within one of the descriptor tables at which the target segment descriptor (having the greater privilege level) is stored. The offset identifies the address at which code fetching is to begin. In 32/64 mode, since the code segment has no base address and flat linear addressing is used, the offset is the address at which code fetching begins. In other modes, the offset is added to the segment base defined by the target segment descriptor to generate the address at which code fetching begins. As mentioned above, the offset may comprise 64 bits in the present embodiment.

DPL 128 stores the minimum privilege level of the calling routine must have (both in the current privilege level and the requested privilege level) which may successfully pass through the call gate and execute the called routine at the privilege level specified in the target segment descriptor.

Type field 130 is coded to a call gate descriptor type. In one embodiment, this type is coded as the 32 bit call gate type defined in the x86 architecture. Alternatively, other encodings may be used. Finally, pseudo-type field 132 is coded to an invalid type (e.g. zero) to ensure that if a segment selector identifying the segment table entry storing the upper half of call gate descriptor 120 is presented, then an exception will be signalled by processor 10.

25

It is noted that the lower half of LDT descriptor 96 may be similar to the 32 bit LDT descriptor and the upper half of LDT descriptor 96 may be similar to the upper half of call gate descriptor 120. Similarly, the lower half of TSS descriptors, trap gate descriptors, and interrupt gate descriptors may be similar to the lower half of the

corresponding 32 bit descriptor and the upper half may be similar to the upper half of call gate descriptor 120.

Turning next to Fig. 10, a block diagram of an instruction format 140 for instructions executed by processor 10 is shown. Other embodiments are possible and contemplated. In the embodiment of Fig. 10, instruction format 140 includes a prefix field 142, an opcode field 144, a mod R/M (register/memory) field 146, an SIB (scale index base) field 148, a displacement field 150, and an immediate field 152. Each of the fields except for the opcode field 144 are optional. Thus, instruction format 140 may define a variable length instruction.

Prefix field 142 is used for any instruction prefixes for the instruction. As described above, an operand size override prefix and an address size override prefix may be encoded into an instruction to override the operating mode of processor 10. These override prefixes are included in prefix field 142. As noted above, the operand size override prefix and address size override prefix may each be bytes included within prefix field 142.

Opcode field 144 includes the opcode of the instruction (i.e. which instruction in the instruction set is being executed). For some instructions, operands may be specified within opcode field 144. For other instructions, a portion of the opcode may be included within mod R/M field 146. Furthermore, certain opcodes specify an eight bit or 16 bit register as an operand. Thus opcode encodings may serve to override the defaults indicated by the operating mode of processor 10 as well.

25

Mod R/M field 146 and SIB field 148 indicate operands of the instruction. Displacement field 150 includes any displacement information, and immediate field 152 includes an immediate operand.

Computer Systems

Turning now to Fig. 11, a block diagram of one embodiment of a computer system 200 including processor 10 coupled to a variety of system components through a bus bridge 202 is shown. Other embodiments are possible and contemplated. In the 5 depicted system, a main memory 204 is coupled to bus bridge 202 through a memory bus 206, and a graphics controller 208 is coupled to bus bridge 202 through an AGP bus 210. Finally, a plurality of PCI devices 212A-212B are coupled to bus bridge 202 through a PCI bus 214. A secondary bus bridge 216 may further be provided to accommodate an electrical interface to one or more EISA or ISA devices 218 through an EISA/ISA bus 10 220. Processor 10 is coupled to bus bridge 202 through a CPU bus 224 and to an optional L2 cache 228. Together, CPU bus 224 and the interface to L2 cache 228 may comprise an external interface to which external interface unit 18 may couple.

Bus bridge 202 provides an interface between processor 10, main memory 204, 15 graphics controller 208, and devices attached to PCI bus 214. When an operation is received from one of the devices connected to bus bridge 202, bus bridge 202 identifies the target of the operation (e.g. a particular device or, in the case of PCI bus 214, that the target is on PCI bus 214). Bus bridge 202 routes the operation to the targeted device. Bus bridge 202 generally translates an operation from the protocol used by the source 20 device or bus to the protocol used by the target device or bus.

In addition to providing an interface to an ISA/EISA bus for PCI bus 214, secondary bus bridge 216 may further incorporate additional functionality, as desired. An input/output controller (not shown), either external from or integrated with secondary bus 25 bridge 216, may also be included within computer system 200 to provide operational support for a keyboard and mouse 222 and for various serial and parallel ports, as desired. An external cache unit (not shown) may further be coupled to CPU bus 224 between processor 10 and bus bridge 202 in other embodiments. Alternatively, the external cache may be coupled to bus bridge 202 and cache control logic for the external cache may be

integrated into bus bridge 202. L2 cache 228 is further shown in a backside configuration to processor 10. It is noted that L2 cache 228 may be separate from processor 10, integrated into a cartridge (e.g. slot 1 or slot A) with processor 10, or even integrated onto a semiconductor substrate with processor 10.

5

Main memory 204 is a memory in which application programs are stored and from which processor 10 primarily executes. A suitable main memory 204 comprises DRAM (Dynamic Random Access Memory). For example, a plurality of banks of SDRAM (Synchronous DRAM) or Rambus DRAM (RDRAM) may be suitable.

10

PCI devices 212A-212B are illustrative of a variety of peripheral devices such as, for example, network interface cards, video accelerators, audio cards, hard or floppy disk drives or drive controllers, SCSI (Small Computer Systems Interface) adapters and telephony cards. Similarly, ISA device 218 is illustrative of various types of peripheral devices, such as a modem, a sound card, and a variety of data acquisition cards such as GPIB or field bus interface cards.

15

Graphics controller 208 is provided to control the rendering of text and images on a display 226. Graphics controller 208 may embody a typical graphics accelerator generally known in the art to render three-dimensional data structures which can be effectively shifted into and from main memory 204. Graphics controller 208 may therefore be a master of AGP bus 210 in that it can request and receive access to a target interface within bus bridge 202 to thereby obtain access to main memory 204. A dedicated graphics bus accommodates rapid retrieval of data from main memory 204. For certain operations, graphics controller 208 may further be configured to generate PCI protocol transactions on AGP bus 210. The AGP interface of bus bridge 202 may thus include functionality to support both AGP protocol transactions as well as PCI protocol target and initiator transactions. Display 226 is any electronic display upon which an image or text can be presented. A suitable display 226 includes a cathode ray tube

("CRT"), a liquid crystal display ("LCD"), etc.

It is noted that, while the AGP, PCI, and ISA or EISA buses have been used as examples in the above description, any bus architectures may be substituted as desired. It
5 is further noted that computer system 200 may be a multiprocessing computer system including additional processors (e.g. processor 10a shown as an optional component of computer system 200). Processor 10a may be similar to processor 10. More particularly, processor 10a may be an identical copy of processor 10. Processor 10a may be connected to bus bridge 202 via an independent bus (as shown in Fig. 11) or may share CPU bus
10 224 with processor 10. Furthermore, processor 10a may be coupled to an optional L2 cache 228a similar to L2 cache 228.

Turning now to Fig. 12, another embodiment of a computer system 300 is shown. Other embodiments are possible and contemplated. In the embodiment of Fig. 12,
15 computer system 300 includes several processing nodes 312A, 312B, 312C, and 312D. Each processing node is coupled to a respective memory 314A-314D via a memory controller 316A-316D included within each respective processing node 312A-312D. Additionally, processing nodes 312A-312D include interface logic used to communicate between the processing nodes 312A-312D. For example, processing node 312A includes
20 interface logic 318A for communicating with processing node 312B, interface logic 318B for communicating with processing node 312C, and a third interface logic 318C for communicating with yet another processing node (not shown). Similarly, processing node 312B includes interface logic 318D, 318E, and 318F; processing node 312C includes interface logic 318G, 318H, and 318I; and processing node 312D includes interface logic 318J, 318K, and 318L. Processing node 312D is coupled to communicate with a plurality of input/output devices (e.g. devices 320A-320B in a daisy chain configuration) via interface logic 318L. Other processing nodes may communicate with other I/O devices in a similar fashion.

30 Processing nodes 312A-312D implement a packet-based link for inter-processing

node communication. In the present embodiment, the link is implemented as sets of unidirectional lines (e.g. lines 324A are used to transmit packets from processing node 312A to processing node 312B and lines 324B are used to transmit packets from processing node 312B to processing node 312A). Other sets of lines 324C-324H are used
5 to transmit packets between other processing nodes as illustrated in Fig. 12. Generally, each set of lines 324 may include one or more data lines, one or more clock lines corresponding to the data lines, and one or more control lines indicating the type of packet being conveyed. The link may be operated in a cache coherent fashion for communication between processing nodes or in a noncoherent fashion for communication
10 between a processing node and an I/O device (or a bus bridge to an I/O bus of conventional construction such as the PCI bus or ISA bus). Furthermore, the link may be operated in a non-coherent fashion using a daisy-chain structure between I/O devices as shown. It is noted that a packet to be transmitted from one processing node to another may pass through one or more intermediate nodes. For example, a packet transmitted by
15 processing node 312A to processing node 312D may pass through either processing node 312B or processing node 312C as shown in Fig. 12. Any suitable routing algorithm may be used. Other embodiments of computer system 300 may include more or fewer processing nodes than the embodiment shown in Fig. 12.

20 Generally, the packets may be transmitted as one or more bit times on the lines 324 between nodes. A bit time may be the rising or falling edge of the clock signal on the corresponding clock lines. The packets may include command packets for initiating transactions, probe packets for maintaining cache coherency, and response packets from responding to probes and commands.

25 Processing nodes 312A-312D, in addition to a memory controller and interface logic, may include one or more processors. Broadly speaking, a processing node comprises at least one processor and may optionally include a memory controller for communicating with a memory and other logic as desired. More particularly, each

processing node 312A-312D may comprise one or more copies of processor 10. External interface unit 18 may include the interface logic 318 within the node, as well as the memory controller 316.

5 Memories 314A-314D may comprise any suitable memory devices. For example, a memory 314A-314D may comprise one or more RAMBUS DRAMs (RDRAMs), synchronous DRAMs (SDRAMs), static RAM, etc. The address space of computer system 300 is divided among memories 314A-314D. Each processing node 312A-312D may include a memory map used to determine which addresses are mapped to which 10 memories 314A-314D, and hence to which processing node 312A-312D a memory request for a particular address should be routed. In one embodiment, the coherency point for an address within computer system 300 is the memory controller 316A-316D coupled to the memory storing bytes corresponding to the address. In other words, the memory controller 316A-316D is responsible for ensuring that each memory access to the 15 corresponding memory 314A-314D occurs in a cache coherent fashion. Memory controllers 316A-316D may comprise control circuitry for interfacing to memories 314A-314D. Additionally, memory controllers 316A-316D may include request queues for queuing memory requests.

20 Generally, interface logic 318A-318L may comprise a variety of buffers for receiving packets from the link and for buffering packets to be transmitted upon the link. Computer system 300 may employ any suitable flow control mechanism for transmitting packets. For example, in one embodiment, each interface logic 318 stores a count of the number of each type of buffer within the receiver at the other end of the link to which that 25 interface logic is connected. The interface logic does not transmit a packet unless the receiving interface logic has a free buffer to store the packet. As a receiving buffer is freed by routing a packet onward, the receiving interface logic transmits a message to the sending interface logic to indicate that the buffer has been freed. Such a mechanism may be referred to as a "coupon-based" system.

I/O devices 320A-320B may be any suitable I/O devices. For example, I/O devices 320A-320B may include network interface cards, video accelerators, audio cards, hard or floppy disk drives or drive controllers, SCSI (Small Computer Systems Interface) adapters and telephony cards, modems, sound cards, and a variety of data acquisition cards such as GPIB or field bus interface cards.

Another Embodiment

Turning next to Fig. 13, a second embodiment of processor 10 is shown. In the embodiment of Fig. 13, control register 26 includes an LMA indication and control register 28 includes a PG indication. Each of these indications may be bits, similar to the LME indication discussion above. The PG bit may indicate whether or not paging is enabled in processor 10. In the embodiment of Fig. 13, processor 10 may not operate as if the LME bit is set until paging is enabled via the PG bit. The LMA bit is used to indicate that both the LME bit is set and the PG bit is set. The LMA bit may be used by processor 10 instead of the LME bit in the above description of Figs. 1-10 to indicate whether or not 32/64 mode is available for use and to indicate compatibility mode for 16 and 32 bit operating modes.

By restricting the use of the processor features enabled by the LME bit to times at which paging is enabled, the concept of a mode in which these features are enabled but paging is not enabled may be avoided. Furthermore, various system tables (such as the page tables) may have a different structure when LME is set and when LME is cleared. Thus, to avoid problems between different page table structures when changing modes in processor 10, the LME bit may be set first, and then the PG bit may be set to enable paging. Since the paging is enabled with the LME bit set, the page table structure defined for the LME bit being set may be used. If paging were enabled prior to the LME bit being set, processor 10 would use the page table structures defined for standard 32 bit mode.

Register Extension (REX) Prefix Byte

One embodiment of processor 10 as illustrated in Figs. 1 and 13 may use a REX prefix byte 400 as illustrated in Fig. 14 to provide for operand size overrides and to provide for additional general registers and additional Streaming Single Instruction,

5 Multiple Data (SIMD) Extension (SSE) registers. REX prefix byte 400 includes an identifier field including bits 7 through 4 of the byte, which identify the prefix byte as the REX prefix byte (e.g. with an encoding of 4 in the illustrated embodiment). Furthermore, the REX prefix byte includes an override bit (OR), and three register specifier bits (R1, R2, and R3).

10

The override (OR) bit may be set to override the default operand size in 32/64 bit mode to 64 bits. If the override bit is clear, the default operand size is used. The operand size override prefix byte (66H) may be used to override the default operand size to 16 bits. Thus, embodiments using the REX byte may use a single override prefix byte to 15 achieve overrides to each desired size, in contrast to the embodiment shown in Fig. 6 for 32/64 mode. The address size override prefix byte may be used similar to the table of Fig. 6.

20

Additionally, the R1, R2, and R3 bits may each be used as the most significant register specifier bit for various register specifiers in the instruction. In other words, one of the R1, R2, and R3 bits may be concatenated with a register specifier drawn from the opcode, Mod R/M, or SIB byte to provide addressing for 16 registers instead of 8 registers. In this manner, the number of registers available for use may be expanded. More particularly, the R1 bit may be concatenated with the register specifier in the "reg" field of the Mod R/M byte. The R2 bit may be concatenated with the register specifier in the index field of the SIB byte. The R3 bit may be concatenated with the register specifier in the "R/M" field of the Mod R/M byte, the base field of the SIB byte, or the "reg" field of the opcode byte. It is noted that a given x86 instruction has a register specifier in only one of the 3 fields specified for the R3 bit.

Software Embodiments

While the above description may generally have described a processor which may directly support, in hardware, the processor architecture having the features described 5 above, it is contemplated that other processor embodiments may not directly implement the processor architecture. Instead, such embodiments may directly implement a different processor architecture (referred to below as a native processor architecture, which may define a native instruction set including native instructions). Any native processor architecture may be used. For example, the MIPS, Power PC, Alpha, Sparc, ARM, etc. 10 architectures may be used. The processor architecture may be implemented in software executing on the native processor architecture in a variety of fashions, using any native processor architecture such as, for example, the Crusoe products of Transmeta Corporation.

15 Generally, a processor embodiment implementing a native processor architecture different than the processor architecture described above (referred to below as the non-native processor architecture) may support the non-native processor architecture in a variety of fashions. For example, such a processor embodiment may execute interpreter software which reads each non-native instruction in a non-native code sequence as data, 20 and executes various software routines which emulate the defined operation of the non-native instruction as defined in the non-native processor architecture. Alternatively, translator software may be executed. The translator software may translate the non-native instructions in the code sequence to an equivalent set of native instructions defined by the native instruction set architecture. The native code sequence may be stored in memory, 25 and may be executed instead of the corresponding non-native code sequence. In yet another alternative, a mixture of interpretation and translation may be used. For example, the code sequence may be interpreted, but the interpreter may also generate statistics about which parts of the code sequence are being most frequently executed. The most frequently executed portions may then be translated to native code sequences.

In any of the above methods, the architected state defined by the non-native processor architecture may be maintained by the combination of the processor and the software (interpreter or translator) in a variety of fashions. For example, the non-native 5 architected state may be mapped to memory locations in a memory addressable by the processor, to general registers defined by the native processor architecture (by software convention, either in the interpreter or in the translator), or the processor may directly support the non-native architected state by defining registers or other storage hardware within the processor that corresponds to the non-native architected state. The non-native 10 architected state may be stored using any combination of the above methods, as desired.

Generally, the architected state includes any state defined to exist by the architecture. For example, in the above described embodiment, the non-native architected state may include general registers (e.g. RAX, RBX, etc.), segment registers, 15 control registers, other registers such as the model specific registers (MSRs), etc. Additionally, the architected state may include data structures defined for the operating system to create, such as the descriptor tables, page tables, task state segments, etc.

Turning to Fig. 15, a flowchart illustrating an exemplary interpreter which may be 20 used to interpret non-native instructions is shown. Other embodiments are possible and contemplated. While the blocks shown are illustrated in a particular order for ease of understanding, any suitable order may be used. Furthermore, blocks may be performed in parallel, as desired.

The blocks shown in Fig. 15 illustrate the emulation of one non-native instruction. 25 Generally, the interpreter may execute the blocks shown in Fig. 15 for each non-native instruction to be executed according to the non-native code sequence to be executed.

The interpreter may determine the operating mode for the non-native instruction (block 1000). As described above, the operating mode may be determined from the LME and/or LMA bits in control register 26 and the L bit and D bit from the code segment descriptor indicated by the CS segment register. The operating mode may be determined anew from the LME, LMA, L bit, and D bit for each non-native instruction, or the resulting operating mode may be stored in a temporary register for access by the interpreter for each non-native instruction. If the resulting operating mode is stored, the interpreter may update the stored operating mode if an instruction modifies the CS segment register or interrupt or exception handling causes the operating mode to change.

10 As mentioned above, the CS segment register and the control register(s) (which are part of the non-native architected state) may actually be memory locations, general registers, or special purpose registers, or any combination thereof.

15 The interpreter may read the current non-native instruction from memory, and may analyze the non-native instruction to determine the operations to be taken to emulate the non-native instruction (block 1002). The interpreter may read the non-native instruction one byte at a time, or may read a suitable set of consecutive bytes and process the bytes. For example, a native processor architecture in which operands are 32 bit may read 32 bits (4 bytes) of the non-native instruction at a time, and then may process the 20 four bytes before reading any additional bytes.

25 Generally, the interpreter software may decode the non-native instruction in a manner analogous to processor 10 decoding the instruction in hardware. Thus, for the illustrated non-native processor architecture, which is compatible with the x86 processor architecture, the analyzing of the non-native instruction includes analyzing any prefix bytes which may precede the opcode byte, analyzing the opcode byte, analyzing the addressing mode (Mod R/M) byte (if present), and analyzing the scale-index-base (SIB) byte (if present). Prefix bytes may override the operating mode, and may also include register specifier bits (e.g. the REX prefix byte). The opcode byte specifies the operation

to be performed, and in some cases may include a register specifier or may implicitly specify an operand (e.g. the stack or the stack pointer). The Mod R/M byte specifies operands (including any displacement or immediate operands which may follow the Mod R/M byte or the SIB byte, if the SIB byte is present) and may include register specifiers.

- 5 Finally, the SIB byte may include register specifiers. From the information gained from analyzing the non-native instruction, the interpreter has the information to emulate the non-native instruction (including operating mode for the non-native instruction, which specifies the operand size and address size of the non-native instruction, operands, the operation to be performed, etc.).

10

If the non-native instruction includes a memory operand (decision block 1004), the interpreter may calculate the effective address of the instruction (block 1006). If the non-native instruction has a memory operand, some of the operands identified in block 1002 may be address operands used to generate the effective address. Thus, the

- 15 interpreter may read the address operands from the non-native architected state and may add them to generate an effective address. The size of the effective address may be determined by the address size for the instruction, as determined at blocks 1000 and 1002. It is noted that the native processor architecture may support an address size which is less than the address size supported by the non-native processor architecture. For example, in
- 20 one exemplary embodiment described above, the virtual address size may be 48 bits in 32/64 mode. The native processor may, for example, support a virtual address size of 32 bits. In such an embodiment, block 1006 may represent a series of calculations in which the least significant bits (e.g. 32 bits) of the virtual address may be calculated, and any carry from the least significant bits may be carried into a calculation of the most
- 25 significant bits of the virtual address.

The interpreter may then perform the operation specified by the non-native instruction (block 1008). If the non-native instruction includes a memory operand as a source operand, the interpreter may read the memory operand from the effective address

calculated at block 1006. Other operands may be read from the non-native architected state. The operation may include an arithmetic operation, a logical operation, a shift, a move to another storage location, etc. The native processor architecture may support an operand size smaller than the operand size of the instruction. In such cases, performing 5 the operation may include multiple calculations on portions of the operand to calculate the result.

The interpreter determines if the non-native instruction resulted in an exception (decision block 1010). Generally, exceptions may occur throughout the execution of the 10 operations specified by the non-native instruction. For example, accessing a source memory operand may result in a page fault before any of the actual instruction operation is performed. During the operations, various architecturally-defined exceptions may also occur. The interpreter may interrupt processing of the non-native instruction upon detecting an exception, and may branch to exception handler instructions (block 1012). 15 The exception handler may be native code or non-native code, as desired. If the non-native processor architecture specifies the update of any architected state when an exception is taken (e.g. various control registers may store the address of the exception causing instruction, the exception reason, etc.), the interpreter may update the non-native architected state as defined.

20 It is noted that the interpreter software is executing on the native processor, and thus is subject to experiencing exceptions as defined in the native processor architecture. These exceptions may generally be different the exceptions detected by the interpreter software, which are exceptions experienced by the non-native code being interpreted 25 according to the non-native processor architecture.

If no exception occurs during emulation of the non-native instruction, the interpreter may update the non-native architected state according to the definition of the non-native instruction (block 1014). Finally, the interpreter may calculate the next non-

native instruction fetch address to fetch the next instruction (block 1016). The next fetch address may be sequential to the current non-native instruction, or may be a different address (e.g. if the current non-native instruction is a taken branch, the next fetch address may be the target address of the branch instruction).

5

Turning now to Fig. 16, a flowchart illustrating a portion of one embodiment of block 1002 from Fig. 15 is shown. Other embodiments are possible and contemplated. While the blocks shown are illustrated in a particular order for ease of understanding, any suitable order may be used. Furthermore, blocks may be performed in parallel, as
10 desired.

The portion of block 1002 illustrated in Fig. 16 illustrates exemplary operation of the interpreter for certain prefix bytes. Other portions, not shown, may handle other prefix bytes or may handle the opcode byte, Mod R/M byte, and SIB byte.

15

More particularly, the interpreter may detect prefix bytes which override the default operand size or address size for the non-native instruction (decision block 1020). Such prefixes may include the address size override prefix byte, the operand size override prefix byte, or the REX byte. If a size override prefix byte is included, the interpreter
20 updates the operating mode for this non-native instruction to reflect the operand size or address size indicated by the prefix byte (block 1022).

25

The interpreter may also detect the inclusion of the REX prefix byte in a non-native instruction (decision block 1024). If the REX prefix byte is included, the interpreter may extract the register specifier bits from the byte for concatenation with the corresponding register specifiers which may be in the opcode, Mod R/M, or SIB bytes (block 1026). For example, in one embodiment, the interpreter may extract each register specifier bit from the REX byte and shift the extracted bit left by 3 bits. The resulting values may each be held in a variable by the interpreter. If there is not REX byte, the

variables may be set to zero. When the register specifiers are encountered in the subsequent bytes, the corresponding variable may be added to the register specifier to provide the complete register specifier for that operand.

5 Turning now to Fig. 17, a flowchart illustrating a portion of one embodiment of block 1008 from Fig. 15 is shown. Other embodiments are possible and contemplated. While the blocks shown are illustrated in a particular order for ease of understanding, any suitable order may be used. Furthermore, blocks may be performed in parallel, as desired.

10

The portion of block 1008 illustrated in Fig. 17 illustrates exemplary operation of the interpreter related to call gates. Other portions, not shown, may handle other portions of the instruction operation.

15

The interpreter determines if the non-native instruction references a call gate (decision block 1070). If the non-native instruction does not reference a call gate, the interpreter continues with other execution tasks for the non-native instruction. If the non-native instruction does reference a call gate, the interpreter reads the descriptor table entry indexed by the segment selector included in the instruction (block 1072). Additionally, 20 the interpreter determines if the LME or LMA bit is set, depending on the embodiment (decision block 1074). If not, then the complete call gate was read via block 1072 and the interpreter may continue by processing the call gate (block 1078). If so, the interpreter reads the next consecutive descriptor table entry (block 1076) to capture the remainder of the call gate. The interpreter may then continue processing the call gate (block 1078).

25

A similar operation may be performed for the LDT descriptor, in response to the LLDT instruction, or for TSS descriptors, interrupt gates, and trap gates.

It is noted that the interpreter may operate in protected mode, using virtual addresses. In other words, the effective address calculated at block 1006 may be a virtual address which is translated by the translation mechanism specified by the non-native processor architecture to a physical address. The processor may include a translation lookaside buffer (TLB) used to cache translations. The processor may either support reload of the TLB from the non-native translation tables (page tables), or may take an exception on a TLB miss to allow software reload of the TLB.

Turning to Fig. 18, a flowchart illustrating an exemplary translator which may be used to translate non-native instructions in the non-native processor architecture to native instructions in the native processor architecture. Other embodiments are possible and contemplated. While the blocks shown are illustrated in a particular order for ease of understanding, any suitable order may be used. Furthermore, blocks may be performed in parallel, as desired.

The blocks shown in Fig. 18 illustrate the translation of one non-native code sequence responsive to a fetch address for the first instruction in the non-native code sequence. The code translator may translate any number of non-native instructions to produce a translated code sequence having native instructions. For example, the translator may translate from the initial non-native instruction to a basic block boundary (i.e. a branch instruction). Alternatively, the translator may speculatively translate two or more basic blocks or may translate up to a maximum number of non-native or resulting native instructions, if desired.

Generally, the translator may maintain a translation cache which stores translated code sequences previously produced by the translator. The translation cache may identify translated code sequences by the fetch address of the first non-native instruction in the corresponding non-native code sequences. Thus, the translator may determine if a translated code sequence corresponding to the fetch address is stored in the translation

cache (decision block 1030). If there is a translated code sequence in the translation cache, the translator may cause the processor to branch to that translated code sequence (block 1032). On the other hand, if there is no translated code sequence, the translator may translate one or more non-native instructions from the non-native code sequence into native instructions in a translated code sequence (block 1034).

Generally, the translator may translate each non-native instruction into one or more native instructions which, when executed, may perform the same operation on the non-native architected state that the non-native instruction would have performed. The translator may generally perform the same decoding of instructions as is performed by the interpreter (block 1002 in Fig. 15) to determine what operations may need to be performed. For example, if the native processor architecture is a load/store architecture in which memory operands are accessed using explicit load/store instructions and other instruction use only register operands, load and store instructions may be used to access the memory operands and other instructions may be used to perform the explicit operation of a non-native instruction having a memory operand. The translated instructions may make use of temporary registers to hold intermediate values corresponding to the execution of the non-native instruction. Additionally, the translated instructions may access the non-native architected state to retrieve operands and may update the non-native architected state with the final results of the non-native instruction. Generally, the native instructions corresponding to the non-native instruction may perform all of the operations defined for the instruction (e.g. blocks 1006, 1008, 1010, 1014, and 1016 in Fig. 15).

Once the translator has determined to terminate translation and save the translated sequence for execution, the translator may optionally optimize the translated code sequence (block 1036). The optimizations may include reordering the translated instructions for quicker execution, eliminating redundancies (e.g. redundant memory references, which may occur if multiple non-native instructions in the source code sequence accessed the same memory location), etc. Any suitable set of optimizations

may be used. The resulting translated code sequence may then be stored into the translation cache. Additionally, the processor may branch to the translated code sequence and execute the sequence (block 1032).

5 It is noted that, while the above description may refer to accessing and/or updating non-native architected state, including various registers, the non-native architected state may be stored in any suitable fashion. For example, architected registers may actually be stored in memory locations, as highlighted above. The mapping of architected registers from the non-native processor architecture to memory locations may be used in either of
10 the interpreter or the translator embodiments, or combinations thereof, to locate the non-architected state used during execution of the non-native instruction or affected by the execution of the non-native instruction. Thus, instructions which access the non-native architected state may perform memory reads/writes or register reads/writes, as the case may be.

15 Turning next to Fig. 19, a block diagram illustrating one exemplary mapping of non-native architected state to either memory locations in a memory 1040 or to processor resources in a native processor 1042. Native processor 1042 includes a register file 1044 including the architected general registers of the native processor architecture. Any
20 number of registers may be provided.

In the embodiment of Fig. 19, all of the non-native architected state is mapped to memory 1040. For example, descriptor tables 1046 (which may include global descriptor table 80, local descriptor table 82, and an interrupt descriptor table), page tables 1048
25 (which store virtual to physical address translations), task state segments 1050, general registers 1052, segment registers 1054, control registers 1056, and other registers 1058 may represent non-native architected state.

Thus, in the embodiment of Fig. 19, to access any non-native architected state, a memory access may be performed. For example, if a non-native instruction has one of the general registers as an operand, the interpreter or translated native instruction performs a memory access to the memory location mapped to that general register to 5 access or update that general register. The registers in register file 1044 may be used by the interpreter or translator as temporary registers to hold intermediate results or for other local interpreter/translator state.

General registers 1052 may include integer general registers (e.g. RAX, RBX, etc. 10 as described above), the additional integer general registers defined by the REX prefix byte, floating point registers, Streaming Single Instruction, Multiple Data (SIMD) Extension (SSE) registers, and the additional SSE registers defined by the REX prefix byte.

15 Segment registers 1054 may include storage locations corresponding to the segment registers 24 shown in Fig. 1 above.

Control registers 1056 may include storage locations corresponding to various control registers defined in the non-native processor architecture. For example, control 20 registers storing the LMA, LME, PG and PE bits, as well as the LDTR and GDTR registers and the CR3 register (which stores the base address of the page tables 1048) are shown. Other control registers may be included as well.

Other registers 1058 includes any remaining architected registers. For example, 25 the EFLAGS register (which stores condition code information), the instruction pointer (RIP) register (which stores the address of the instruction to be executed), and the model specific registers (MSRs) may be included in other registers 1058.

While the example of Fig. 19 maps all of the non-native architected state to memory 1040, other embodiments may implement other mappings. In Fig. 20, for example, some of the general registers in register file 1044 are mapped to the general registers 1052. Accordingly, if a non-native instruction has a general register as an operand, the interpreter accesses the corresponding register in register file 1044. Similarly, the translator generates a translated instruction having the corresponding register in register file 1044 as an operand. Other architected state may still be accessed via memory operations in the embodiment of Fig. 20. Other registers in register file 1044 which are not assigned to non-native architected state may again be used as temporary registers for interpreter or translator use, as described above.

While the embodiment of Fig. 20 illustrates mapping the general registers 1052 to registers in register file 1044, any other non-native architected state may be mapped to registers in register file 1044. For example, any of segment registers 1054, control registers 1056, or other registers 1058 (or portions of any of these registers) may be mapped to register file 1044, as desired.

Fig. 21 illustrates another example in which the general registers 1052 and the EFLAGS and RIP registers are mapped to registers in register file 1044. Additionally, in the example of Fig. 21, the segment registers 1054 are implemented in hardware in processor 1042. More specifically, processor 1042 may not only implement storage for segment registers 1054, but may also include logic to generate the operating mode for instructions based on the information in the segment registers. Furthermore, for compatibility modes, the logic may include limit checks and attribute checks to ensure that accesses to the segment attempted by the non-native instructions (or the non-native instructions in the interpreter or the translated code sequence which correspond to the non-native instructions) are permitted.

Similarly, other embodiments may implement various control registers 1056 or other registers 1058 in hardware, including corresponding logic to act on the contents of the registers as defined in the non-native architecture. Generally, various embodiments of processor 1042 may implement any non-native architected state in hardware. Certain
5 architected state may generally be implemented in memory since the non-native processor architecture defines the state to be in memory (e.g. descriptor tables 1046, pages tables 1048, and task state segments 1050). Such memory-based architected state may be cached in caches within processor 1042 (e.g. TLBs for page table information, hidden segment register portions for segment descriptor information, etc.).

10

As the above discussion illustrates, the non-native architected state may be stored in any suitable storage location. Generally, a storage location is a location capable of storing a value. Suitable storage locations may include, in various embodiments, a memory location, a general register mapped to the non-native architected state, or a
15 special purpose register (which may include additional hardware to interpret the contents of the register), depending upon the embodiment. Additionally, suitable storage locations could include a scratch pad RAM (such as a portion of a cache predetermined to be used as scratch pad RAM).

20

Fig. 22 is a block diagram of one embodiment of a carrier medium 1090. Other embodiments are possible and contemplated. In the embodiment of Fig. 22, carrier medium 1090 stores an interpreter program 1092 and a translator program 1094.

25

Generally speaking, a carrier medium may include storage media such as magnetic or optical media, e.g., disk or CD-ROM, volatile or non-volatile memory media such as RAM (e.g. SDRAM, RDRAM, SRAM, etc.), ROM, etc., as well as transmission media or signals such as electrical, electromagnetic, or digital signals, conveyed via a communication medium such as a network and/or a wireless link. Carrier medium 1090 may thus be coupled to a computer system including processor 1042, may be part of a

computer system including processor 1042, or may be a communication medium on which the computer system is capable of communicating. Computer systems including processor 1042 may be of any construction. For example, computer systems similar to those shown in Figs. 11 and 12 may be suitable.

5

Interpreter program 1090 may operate according to the flowchart of Fig. 15 (and subsequent flowcharts illustrating portions of Fig. 15 in greater detail). Translator program 1094 may operate according to the flowchart of Fig. 18. Generally, interpreter program 1092 and translator program 1094 may each comprise code sequences including
10 native instructions.

Numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

15